

Glance-or-Gaze: Incentivizing LMMs to Adaptively Focus Search via Reinforcement Learning

Hongbo Bai^{1*}, Yujin Zhou^{1*}, Yile Wu¹, Chi-Min Chan¹
Pengcheng Wen¹, Kunhao Pan¹, Sirui Han^{1†}, Yike Guo^{1†}

¹Hong Kong University of Science and Technology
baihongbo@ust.hk
yzhouha@connect.ust.hk

Abstract

Large Multimodal Models (LMMs) have achieved remarkable success in visual understanding, yet they struggle with knowledge-intensive queries involving long-tail entities or evolving information due to static parametric knowledge. Recent search-augmented approaches attempt to address this limitation, but existing methods rely on indiscriminate whole-image retrieval that introduces substantial visual redundancy and noise, and lack deep iterative reflection, limiting their effectiveness on complex visual queries. To overcome these challenges, we propose Glance-or-Gaze (GoG), a fully autonomous framework that shifts from passive perception to active visual planning. GoG introduces a Selective Gaze mechanism that dynamically chooses whether to glance at global context or gaze into high-value regions, filtering irrelevant information before retrieval. We design a dual-stage training strategy: Reflective GoG Behavior Alignment via supervised fine-tuning instills the fundamental GoG paradigm, while Complexity-Adaptive Reinforcement Learning further enhances the model’s capability to handle complex queries through iterative reasoning. Experiments across six benchmarks demonstrate state-of-the-art performance. Ablation studies confirm that both Selective Gaze and complexity-adaptive RL are essential for effective visual search. We will release our code and models for further exploration soon.¹

“You see, but you do not observe. The distinction is clear.”

— Arthur Conan Doyle, *A Scandal in Bohemia*

1 Introduction

Large Multimodal Models (LMMs) have demonstrated remarkable proficiency in general visual

understanding and reasoning, driven by extensive visual-text pre-training (Wang et al., 2025b, 2024; Qi et al., 2025). Despite these progress, a fundamental tension persists between the static nature of parametric knowledge and the dynamic, open-ended evolution of the real world. Constrained by fixed training corpora, LMMs inevitably struggle with knowledge-intensive Visual Question Answering (VQA) tasks that require up-to-date information or involve concepts of long-tail beyond their internal distribution (Narayan et al., 2025; Li et al., 2023). Specifically, when querying recent events or obscure visual entities, such as identifying a niche festival or interpreting a newly released product, these models frequently exhibit “knowledge cutoff” or “long-tail amnesia” (Lewis et al., 2020; Hu et al., 2023; Xu et al., 2024; Tong et al., 2024). Such deficiencies often precipitate factual hallucinations or generic responses, thereby undermining the LMMs’ utility in knowledge-critical domains (Wei et al., 2024; Yasunaga et al., 2022; Luo et al., 2025; Ji et al., 2025).

To mitigate this gap, recent research has increasingly sought to augment LMMs with external search mechanisms. Existing works generally fall into three distinct categories. Initially, approaches adopt Retrieval-Augmented Generation (RAG), which retrieves auxiliary context from fixed knowledge bases (Wu et al., 2025b; Lin et al., 2024, 2023). However, this “retrieve-then-generate” pipeline often suffers from rigidity, failing to capture the dynamic breadth of the open web. Subsequently, Prompt-Engineered Agents emerged, leveraging the planning capabilities of LMMs to orchestrate search engines (Li et al., 2025; Wu and Xie, 2024). While flexible, these systems typically operate in a “plug-and-play” manner, leaving the model’s intrinsic search capabilities unoptimized. More recently, tool-integrated LMMs have advanced this frontier by incorporating search actions directly into the training process,

*Equal contribution. †Corresponding author.

¹<https://github.com/TOM-ZHOUC/Glance-or-Gaze>

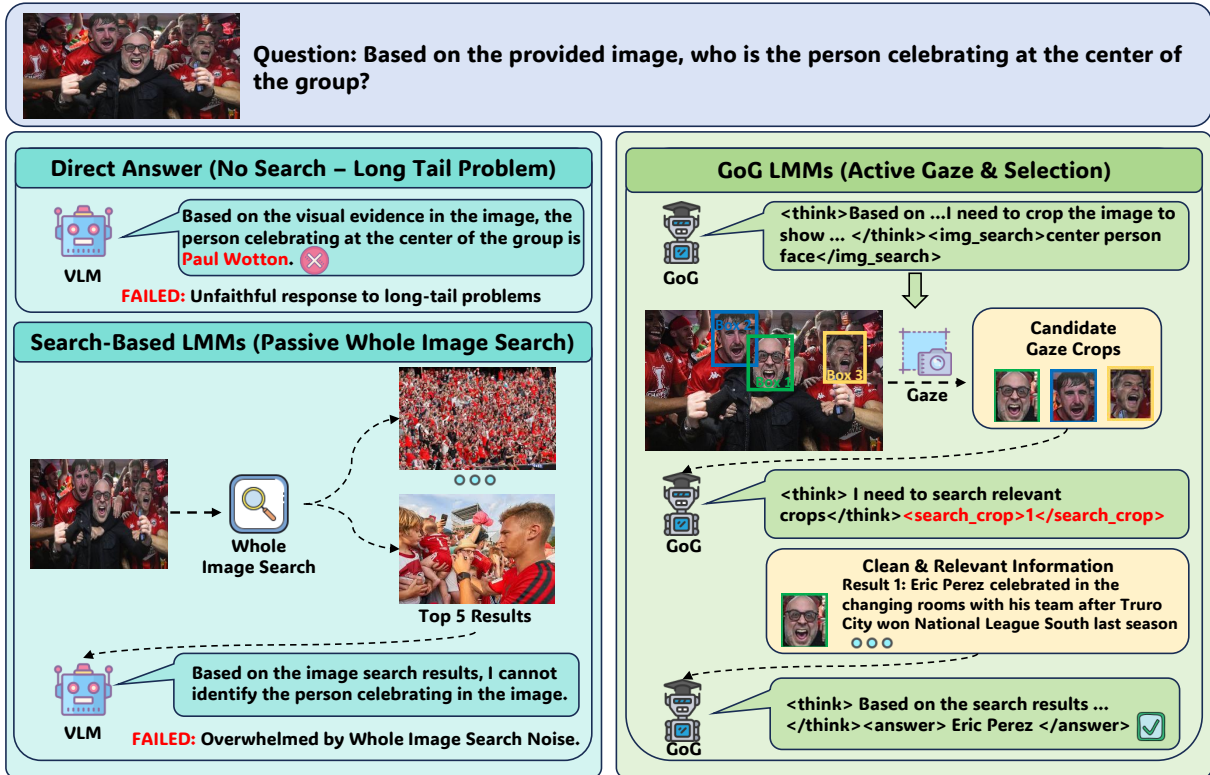


Figure 1: Comparison between previous baselines and our Glance-or-Gaze (GoG) framework. GoG employs an active multi-step strategy: proposing candidate regions via visual grounding, filtering relevant crops through Selective Gaze, and conducting precise search only on selected regions with iterative cross-modal reflection.

thereby aligning the model’s internal representations with external retrieval behaviors (Wu et al., 2025a; Narayan et al., 2025).

Despite recent progress, existing methodologies face two critical limitations. First, current visual strategies are inefficient and noisy. By relying on rigid cropping and indiscriminately processing all candidate regions, these methods introduce significant noise. Moreover, they depend heavily on converting visual details into text descriptions, which creates an information bottleneck and prevents the model from accessing raw visual data. Second, current works execute tool calls in a single-pass manner or limit reflection strictly to the textual modality. They fail to support iterative visual verification or self-correction, restricting their ability to adaptively refine strategies when facing intricate visual questions.

To overcome these limitations, we propose **Glance-or-Gaze (GoG)**, the first fully autonomous framework that shifts the paradigm from passive image perception to dynamic, complexity-adaptive visual planning, as illustrated in Figure 1. We introduce a Selective Gaze mechanism that actively filters visual noise by evaluating and prioritizing

pertinent image patches before execution. The framework integrates two complementary learning mechanisms: (i) **Reflective GoG Behavior Alignment** addresses the inefficiency and noise inherent in current visual strategies by employing imitation learning to instill the fundamental paradigm of active selection and cross-modal reflection. (ii) **Complexity-Adaptive Reinforcement Learning** further enhances the model’s planning and reasoning capabilities by optimizing the selection policy, enabling intelligent, iterative reflection and adaptive focus search based on visual query complexity. Our main contributions are summarized as follows:

- We propose GoG, the first fully autonomous framework that shifts the paradigm from passive image perception to dynamic, complexity-adaptive visual planning. We introduce a Selective Gaze mechanism that actively filters visual noise by evaluating and prioritizing pertinent image patches, effectively bridging the gap between coarse-grained glancing and fine-grained evidence verification.
- We design a dual-stage learning architecture comprising Reflective GoG Behavior Align-

ment and Complexity-Adaptive Reinforcement Learning. The former addresses the inefficiency inherent in current visual strategies by instilling the fundamental paradigm of active selection and cross-modal reflection. The latter further enhances planning and reasoning capabilities, enabling adaptive search based on visual query complexity.

- Extensive experiments demonstrate the effectiveness and generalization of GoG. Our method achieves state-of-the-art performance across diverse benchmarks, surpassing strong baselines with significant gains ranging from 5 to 20, validating the superiority of our adaptive GoG planning paradigm.

2 Related Works

2.1 Large Multimodal Models

The evolution of Large Multimodal Models (LMMs) has progressed rapidly from early vision-language pretraining (Agarwal et al., 2021; Li et al., 2022; Zhu et al., 2025) to sophisticated unified architectures capable of complex reasoning (Liang et al., 2026; Bai et al., 2025b; Yakun et al., 2026b,a; Zhou et al., 2026). Recent models such as GPT-4o (OpenAI et al., 2024), Gemini (Team et al., 2025), Qwen3-VL series (Bai et al., 2025a), and Internvl3.5-VL (Wang et al., 2025a) have demonstrated remarkable capabilities in visual understanding, achieving strong performance on benchmarks spanning image captioning and visual question answering. Despite these advances, LMMs remain fundamentally constrained by their static parametric knowledge (Lewis et al., 2020; Hu et al., 2023; Liu et al., 2026). When confronted with queries involving recent events, long-tail entities, or fine-grained visual details beyond their training distribution, these models frequently hallucinate or produce generic responses (Xu et al., 2024).

2.2 Search-Augmented Visual Reasoning

To address the knowledge limitations of LMMs, researchers have explored various strategies for integrating external information retrieval. Early efforts extended text-based Retrieval-Augmented Generation (RAG) to the multimodal domain, retrieving relevant documents or images from fixed knowledge bases (Wu et al., 2025b; Lin et al., 2024, 2023). Subsequently, prompt-based search agents emerged, leveraging the planning capabilities of

LMMs to orchestrate web search engines via chain-of-thought prompting (Yao et al., 2022; Yang et al., 2023; Li et al., 2025; Wu and Xie, 2024). More recently, tool-integrated LMMs have incorporated search actions directly into the training process, with MMSearch-R1 (Wu et al., 2025a) embedding search within the reinforcement learning loop and optimizing web navigation policies through online interaction (Wu et al., 2025a; Narayan et al., 2025). Despite these advances, current visual strategies remain inefficient and noisy, relying on rigid cropping and indiscriminate region processing while converting visual details into text, creating an information bottleneck that prevents access to raw visual data. Furthermore, existing approaches execute tool calls in a single-pass manner without iterative visual verification, limiting their ability to adaptively refine strategies for complex visual queries. To address these challenges, we propose Glance-or-Gaze (GoG), a unified framework that integrates active visual exploration with reflective reasoning, enabling dynamic, complexity-adaptive visual planning.

3 Method

3.1 Overview

We propose a dual-stage learning framework designed to evolve Large Multimodal Models (LMMs) from passive observers into active visual planners. As illustrated in Figure 2, our approach integrates *Reflective GoG Behavior Alignment* and *Complexity-Adaptive Reinforcement Learning* to achieve complexity-adaptive visual reasoning.

- **Stage 1: Reflective GoG Behavior Alignment.** This stage addresses the inefficiency and noise inherent in current visual strategies. We establish a specialized data curation pipeline with uncertainty-aware query sampling and human-verified trajectory synthesis to construct the **GoG-Instruct** dataset. Through supervised fine-tuning, we instill the fundamental paradigm of active selection and cross-modal reflection, training the model to suppress ungrounded generation and proactively trigger “Glance” or “Gaze” actions.
- **Stage 2: Complexity-Adaptive Reinforcement Learning.** This stage further enhances the model’s planning and reasoning capabilities. We construct a complexity-stratified dataset based on failure rates and

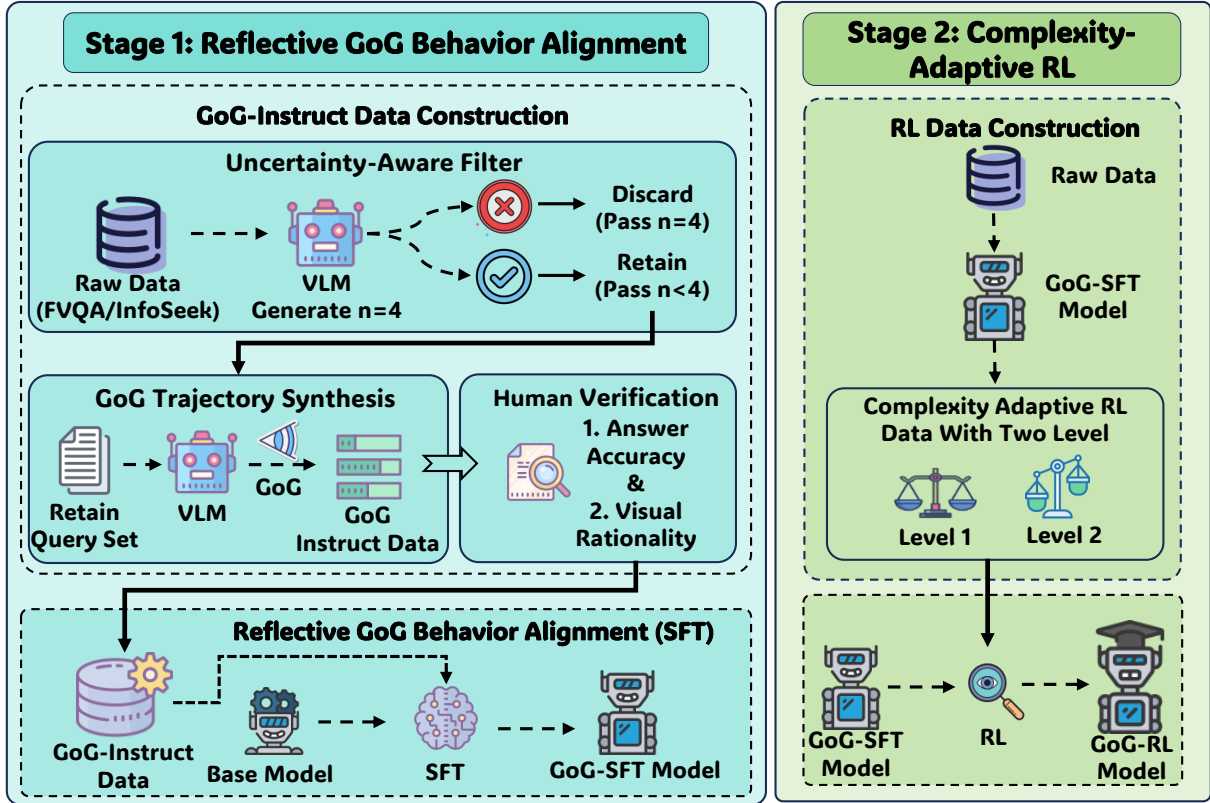


Figure 2: Overview of the Glance-or-Gaze (GoG) framework. **Stage 1 (Left):** Reflective GoG Behavior Alignment constructs GoG-Instruct data through uncertainty-aware filtering and human-verified trajectory synthesis, then performs supervised fine-tuning to instill active selection and cross-modal reflection. **Stage 2 (Right):** Complexity-Adaptive RL constructs complexity-stratified data at two difficulty levels and applies reinforcement learning to enhance planning capabilities for adaptive visual reasoning.

employ Group Relative Policy Optimization (GRPO) (Shao et al., 2024). This empowers the model to autonomously plan and execute multi-step exploratory searches, enabling adaptive search based on visual query complexity.

3.2 Stage 1: Reflective GoG Behavior Alignment

This stage instills the fundamental paradigm of active selection and cross-modal reflection through imitation learning.

GoG-Instruct Data Construction To facilitate active visual planning, We curate *GoG-Instruct*, a dataset derived from FVQA (Wu et al., 2025a) and InfoSeek (Chen et al., 2023), through a three-step pipeline as illustrated in Figure 1 (Left).

1. Uncertainty-Aware Filtering. We filter out trivial samples where the model already possesses sufficient parametric knowledge. Specifically, we employ Qwen3-VL-235B-Instruct to generate answers for each query $N = 4$ times. Samples where the

model consistently answers correctly (pass count $n = 4$) are discarded, retaining only queries with $n < 4$ that require external visual verification.

2. GoG Trajectory Synthesis. For the retained query set, we synthesize reasoning trajectories with explicit visual planning structure. Each trajectory follows the sequence: (1) *Glance* for global image analysis, (2) *Decision* for tool invocation or region cropping, and (3) *Gaze* for targeted search execution on selected regions.

3. Human Verification. To ensure data quality, expert annotators validate each trajectory against two criteria: (1) **Answer Accuracy**, ensuring the final response is factually correct; and (2) **Visual Rationality**, verifying that cropped regions are logically relevant to the query. Trajectories failing either criterion are discarded.

Through this pipeline, we construct the GoG-Instruct dataset comprising 5,750 samples, where 43.5% are search-free instances and 56.5% require various search operations, as detailed in Table 1.

Type	Subtype	Ratio	Count
Search-free	–	43.5%	2,500
Search-required	Text Search Only	13.0%	750
	Image Search Only	30.4%	1,750
	Both Search	13.0%	750
Total		100%	5,750

Table 1: Composition of the SFT training dataset. “Image Search Only” refers to instances that execute either a single whole image search operation or a Gaze operation (cropped region search), but not both.

Training Objective We employ Qwen2.5-VL-7B-Instruct and Qwen3-VL-Think as the base model. To enable efficient training, we incorporate LoRA adapters ($r = 8$) across all transformer blocks. The model is trained on multi-turn conversations y^* containing reasoning sequences and structured tool calls. We adopt the standard Causal Language Modeling (Causal LM) objective:

$$\mathcal{L}_{\text{SFT}} = - \sum_{t=1}^T \log \pi_{\theta}(y_t^* | x, I, y_{<t}^*), \quad (1)$$

where T is the sequence length and π_{θ} is the model’s conditional distribution.

3.3 Stage 2: Complexity-Adaptive Reinforcement Learning

While SFT establishes the capability for active planning, it lacks the flexibility to handle varying query difficulties. We employ Reinforcement Learning to optimize the planning policy for robustness, as illustrated in Figure 1 (Right).

RL Data Construction We construct complexity-stratified data by evaluating the GoG-SFT model on raw queries and stratifying them into two difficulty levels:

Level 1: This subset consists of queries where the GoG-SFT model exhibits a pass rate of approximately 50% using standard “Glance” or “Gaze” actions. These samples represent the decision boundary where the model is uncertain, providing effective signal for policy optimization.

Level 2: Building upon Level 1, this subset incorporates additional queries where the GoG paradigm consistently fails, resulting in an overall lower pass rate. Training on this level forces the model to develop more robust reasoning strategies beyond the standard visual exploration patterns.

We adopt Level 2 data as our final RL training set to maximize the model’s capability for complex

visual reasoning.

Training Objective We employ GRPO to optimize the policy on the Level 2 dataset \mathcal{D}_{L2} . For each multimodal input $(q, I) \in \mathcal{D}_{\text{L2}}$, we sample a group of G candidate trajectories $\{\tau_i\}_{i=1}^G$ from the current policy π_{θ} . Each trajectory receives a reward r_i , and we compute the group-normalized advantage $\hat{A}_i = (r_i - \mu_r) / \sigma_r$, where μ_r and σ_r are the mean and standard deviation of rewards within the group. The objective function is:

$$\mathcal{L}_{\text{GRPO}}(\theta) = \mathbb{E}_{(q,I) \sim \mathcal{D}_{\text{L2}}} \left[\frac{1}{G} \sum_{i=1}^G \mathcal{L}_i^{\text{clip}} - \beta \mathbb{D}_{\text{KL}}(\pi_{\theta} \| \pi_{\text{ref}}) \right] \quad (2)$$

where $\mathcal{L}_i^{\text{clip}}$ is the clipped surrogate objective and β controls the KL divergence penalty to prevent excessive deviation from the reference model π_{ref} .

The reward r_i combines accuracy and format compliance. The accuracy score $r_{\text{acc}} \in \{0, 1\}$ is judged by gpt-oss-120b based on semantic correctness against ground truth. The format score $r_{\text{fmt}} \in [0, 1]$ validates correct usage of special tokens and structural integrity. The total reward is:

$$r_i = (1 - \lambda) \cdot r_{\text{acc}} + \lambda \cdot r_{\text{fmt}} \quad (3)$$

where λ controls the emphasis on formatting compliance.

4 Experiments

4.1 Experiment Setup

Implementation Details We initialize our backbone using Qwen2.5-VL-7B-Instruct and Qwen3-VL-8B-Think. For SFT, we utilize the LLaMA-Factory framework with LoRA (rank 8) applied to all target modules. The models are trained for 3 epochs with a learning rate of $1e^{-4}$, a cosine learning rate scheduler, and a warmup ratio of 0.1. We employ bf16 mixed precision with a global batch size of 8. For online RL optimization, we adopt the Group Relative Policy Optimization (GRPO) algorithm within the veRL framework. We set the actor learning rate to $2e^{-6}$ with a sigmoid decay warmup over 45 steps. The KL-divergence coefficient is set to $\beta = 0.001$. During exploration, we use a rollout number of $N = 4$ with vLLM for efficient inference, allowing a maximum response length of 8,192 tokens. The RL training runs for 15 epochs. All training are conducted on a single node equipped with 8 NVIDIA H800 GPUs. Additional hyperparameter details are provided in Appendix D.

Benchmark and Metrics Benchmarks. To comprehensively evaluate GoG, we conduct experiments on a diverse set of datasets categorized into in-distribution (IID) and out-of-distribution (OOD) settings. For IID evaluation, we utilize InfoSeek (Chen et al., 2023) and FVQA-test (Wu et al., 2025a). To assess the model’s generalization capability (OOD), we employ SimpleVQA (Cheng et al., 2025), MMSearch (Jiang et al., 2024), DynVQA (Li et al., 2024), and LiveVQA-New (Fu et al., 2025). Due to the substantial size of the original splits, we randomly sample 2,000 instances from InfoSeek and LiveVQA for evaluation. Additionally, we filter all datasets to exclusively retain English-language samples. Following prior work, we employ an LLM-as-a-Judge approach to evaluate model performance. Specifically, we utilize gpt-oss-120b as the judge model to directly assess the accuracy of the generated answers. The full evaluation prompt used for this assessment is provided in Appendix E.

Baselines. To benchmark the capabilities of GoG, we evaluate it against a series of strong baselines, including both open-source models (the Qwen2.5-VL and Qwen3-VL series) and closed-source models (GPT-4o (OpenAI et al., 2024)). We primarily compare performance across four distinct settings: (1) **Direct Answer**, where the model is provided with the image and question and asked to respond directly; (2) **Full-Search Workflow**, where retrieval is mandatory for every query; (3) **Prompt-based GoG Agents**, which utilize Graph-of-Thought reasoning via prompting; and (4) **Search-Equipped LMMs**, specifically comparing against MMSearch R1. Detailed evaluation protocols and the specific prompts used are provided in Appendix E.

4.2 Main Results and Observations

GoG achieves state-of-the-art performance with robust generalization across diverse benchmarks. As detailed in Table 2, GoG-3-8B-Think-RL surpasses all baseline paradigms, outperforming the strongest Full-Search Workflow and Prompt-based GoG agent by +9.89 and +15.18 on average, respectively. Notably, compared to the previous search-equipped baseline MMSearch-R1, GoG achieves a remarkable +19.97 improvement. This superiority is consistent across both in-domain datasets and out-of-domain settings, demonstrating the model’s exceptional adaptability. We observe

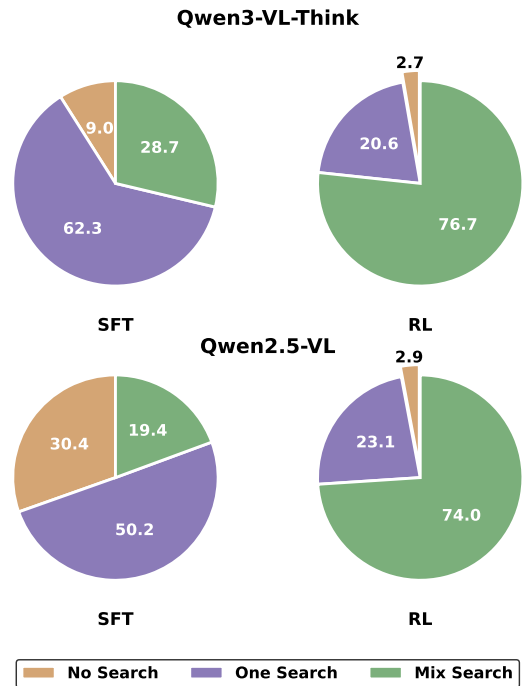


Figure 3: Distribution of search behavior across different training stages. “No Search” indicates samples without any search action, “One Search” represents samples using only one type of search (text, image, or crop), and “Mix Search” denotes samples combining multiple search types.

that while Full-Search workflows improve over Direct Answer baselines (such as GPT-4o) by retrieving external knowledge, they often fall short of GoG because they lack the learned discrimination to filter irrelevant noise or decide when search is strictly necessary.

SFT teaches basic GoG capability while RL enhances interactive GoG reasoning. As shown in Figure 3, SFT and RL exhibit distinct search behavior patterns. After SFT, both models primarily rely on single-type searches (62.3% for Qwen3-VL-Think and 50.2% for Qwen2.5-VL), with a notable proportion of samples requiring no search at all (9.0% and 30.4%, respectively). This suggests that SFT successfully teaches the model *when* and *how* to invoke search tools for straightforward queries. However, after RL training, the proportion of mix search increases dramatically—from 28.7% to 76.7% for Qwen3-VL-Think and from 19.4% to 74.0% for Qwen2.5-VL. Meanwhile, the no-search ratio drops to below 3% for both models. This shift indicates that RL encourages the model to engage in more complex, multi-step information gathering, combining text search, image search, and image cropping to thoroughly verify and cross-reference

Model	Avg.	In-Domain		Out-of-Domain			
		FVQA	InfoSeek	SimpleVQA	MMSearch	LiveVQA	DynVQA
<i>Direct Answer</i>							
Qwen2.5-VL-7B	20.31	25.33	13.90	41.26	14.62	10.30	16.43
Qwen2.5-VL-32B	22.79	26.89	16.95	42.84	16.37	12.70	20.96
Qwen3-VL-8B-Instruct	21.13	23.94	13.80	40.18	15.20	11.40	22.24
Qwen3-VL-8B-Thinking	23.84	24.56	16.05	41.76	15.20	15.15	30.31
Qwen3-VL-30B-A3B-Instruct	23.80	26.11	16.70	42.35	18.13	12.60	26.91
Qwen3-VL-30B-A3B-Thinking	27.26	31.11	22.80	43.73	17.54	15.40	33.00
GPT-4o	30.68	42.00	30.60	43.44	21.64	14.80	31.59
<i>Prompt-based GoG Method</i>							
Qwen2.5-VL-7B	22.31	24.50	17.40	41.66	20.47	10.40	19.41
Qwen3-VL-8B-Think	41.70	51.33	32.00	62.69	36.84	25.95	41.36
<i>Full-Search Workflow</i>							
Qwen2.5-VL-7B	44.36	61.06	38.15	59.13	36.26	31.05	40.51
Qwen3-VL-8B-Thinking	46.99	57.33	32.25	61.90	63.84	23.55	39.09
<i>Search-Equipped Models</i>							
MMSearch-R1*	36.91	42.39	24.65	54.79	40.94	23.85	34.84
GoG-2.5-7B-SFT	43.28	53.72	41.90	60.61	40.35	24.55	38.53
GoG-3-8B-Think-SFT	50.17	62.17	40.55	65.65	53.80	32.40	46.46
GoG-2.5-7B-RL	53.22	66.78	51.05	64.86	56.73	37.70	42.21
GoG-3-8B-Think-RL	56.88	68.44	49.05	66.44	65.50	43.85	48.02

Table 2: Main evaluation results. We compare Direct Answer, Prompt-based GoG, Full-Search Workflow, and Search-Equipped Models. Blue rows indicate our proposed models (SFT and RL). Gray row indicates the reproduced MMSearch-R1 baseline. * denotes reproduced results.




Figure 4: **Effectiveness of Selective Gaze.** Comparison of Crop Selection Accuracy between SFT and RL training stages across two model architectures.

information before answering. The emergence of such iterative search behavior demonstrates that RL not only reinforces correct search usage but also cultivates a more deliberate reasoning process.

RL training sharpens Selective Gaze to identify answer-relevant regions more precisely. Beyond overall performance gains, we investigate whether RL training specifically improves the model’s ability to select relevant image regions. As shown in Figure 4, we measure *Crop Selection Accu-*

racy—the proportion of selected crops that contain answer-relevant information—on our evaluation set. Both models exhibit substantial improvements after RL training: Qwen3-VL-Think improves from 42.1% to 48.9% (+6.8%), while Qwen2.5-VL increases from 46.7% to 51.3% (+4.7%). This consistent gain across different architectures demonstrates that RL does not merely encourage more frequent tool use, but fundamentally refines *where* the model chooses to focus. To gain deeper insights, we manually examined 100 samples from GoG-3-8B-Think where Gaze operations were invoked (Table 3). We evaluated two aspects: (1) *Gaze Correctness*—whether the selected crop is appropriate for answering the question, and (2) *Reflection Rate*—among incorrect selections, whether the model triggers self-correction. Results show that RL training improves Gaze correctness from 59% to 75%, and more strikingly, enhances error awareness: when the initial Gaze is incorrect, SFT models rarely recognize the error (30% reflection rate), whereas RL-trained models trigger reflection 70% of the time. This indicates that RL instills better-calibrated uncertainty, enabling the model to recognize misguided focus and actively seek alternative regions.

Metric	SFT	RL
Gaze Correctness (%)	59	75
Reflection Rate on Errors (%)	30	70

Table 3: Manual analysis of Gaze behavior on 100 GoG-3-8B-Think samples.

Model Setting	Avg.	FVQA	Info	Simple	MM	Live	Dyn
<i>Base Model: Qwen2.5-VL-7B</i>							
SFT w/o SG	41.76	53.39	40.10	59.53	40.94	22.30	34.28
Full SFT	43.28	53.72	41.90	60.61	40.35	24.55	38.53
Δ Gain	+1.52	+0.33	+1.80	+1.08	-0.59	+2.25	+4.25
<i>Base Model: Qwen3-VL-8B-Thinking</i>							
SFT w/o SG	48.34	60.17	40.55	64.86	46.20	32.80	45.47
Full SFT	50.17	62.17	40.55	65.65	53.80	32.40	46.46
Δ Gain	+1.83	+2.00	+0.00	+0.79	+7.60	-0.40	+0.99

Table 4: Ablation study results. We compare the full SFT model against the ablated version (SFT w/o SG), where **SG** denotes the **Selective Gaze mechanism**. Green highlights notable gains; red indicates slight regressions.

4.3 Ablation and Analysis

Effectiveness of the Selective Gaze Mechanism.

To validate our architectural design, we conduct an ablation study by removing the Selective Gaze (SG) mechanism. In this "w/o SG" variant, the model loses the ability to perform dynamic visual cropping and reflection, regressing to a baseline that relies solely on a coarse-grained global view of retrieved images. As shown in Table 4, this removal leads to consistent performance degradation across both Qwen2.5 and Qwen3 backbones, with average scores dropping by 1.52 and 1.83 points, respectively. The results underscore the necessity of the "Selective-Gaze-and-Reflect" paradigm, particularly in complex scenarios like DynVQA (+4.25 gain) and MMSearch (+7.60 gain) where target information is often minute or obscured. The SG mechanism acts as a physical anchor for reasoning: by explicitly directing its gaze to a selected candidate region, it forces the model to transition from passively *seeing* the global context to actively *observing* specific visual details. This focused verification provides the granular evidence necessary to ground the model’s reasoning, preventing hallucinations common in noise-heavy environments.

Effectiveness of Complexity-Adaptive RL Training. We further investigate the impact of data difficulty on policy optimization, positing that samples challenging for the SFT model provide richer learning signals. To this end, we partition the RL

Model Setting	Avg.	FVQA	Info	Simple	MM	Live	Dyn
<i>Base Model: Qwen2.5-VL-7B</i>							
RL w/ Level 1 Data	47.28	64.00	50.35	64.66	50.88	32.95	39.94
RL w/ Level 2 Data	53.22	66.78	51.05	64.86	56.73	37.70	42.21
Δ Gain	+5.94	+2.78	+0.70	+0.20	+5.85	+4.75	+2.27
<i>Base Model: Qwen3-VL-8B-Thinking</i>							
RL w/ Level 1 Data	48.89	66.39	44.95	64.86	61.99	39.55	45.61
RL w/ Level 2 Data	52.38	68.44	49.05	66.44	65.50	43.85	48.02
Δ Gain	+3.49	+2.05	+4.10	+1.58	+3.51	+4.30	+2.41

Table 5: Ablation study on Complexity-Adaptive RL data construction. **Level 2 Data** refers to samples with SFT model pass rate < 50%, while **Level 2 Data** includes samples with pass rate \geq 50%. Green highlights substantial gains (\geq 4%). Training on hard samples consistently yields improvements across both backbones.

training pool into **Level 2** (pass rate < 50%) and **Level 1** (pass rate 50%). As shown in Table 5, training on hard samples consistently outperforms the easy data setting. Specifically, for Qwen2.5-VL-7B, the hard data setting achieves a +5.94 average gain over the easy setting, with significant boosts on benchmarks requiring complex reasoning, such as MMSearch (+5.85) and LiveVQA (+4.75). Similarly, Qwen3-VL-8B-Think demonstrates a +3.49 average improvement with consistent gains across all metrics. These results confirm that while easy samples offer limited room for policy refinement, hard samples present genuine decision-making challenges that maximize the benefits of exploratory optimization, thereby validating our complexity-adaptive data construction strategy.

5 Conclusion

We present Glance-or-Gaze (GoG), a framework that transforms Large Multimodal Models from passive observers into active visual planners. GoG introduces the Selective Gaze mechanism to dynamically allocate attention between global context and fine-grained regions, filtering visual noise before retrieval. Our dual-stage training strategy—Reflective Behavior Alignment via SFT followed by Complexity-Adaptive Reinforcement Learning—progressively instills cross-modal planning capabilities and refines decision-making on challenging queries. Experiments across six benchmarks demonstrate state-of-the-art performance, with GoG surpassing strong baselines by substantial margins. We believe this work establishes a promising direction for building more capable and autonomous multimodal search agents.

6 Limitations

While GoG demonstrates strong performance across diverse benchmarks, several aspects warrant further investigation. First, although we have optimized the stability of our search infrastructure and Jina Reader pipeline, we observe an occasional failure rate of approximately 1–5% due to network instability, API timeouts, or malformed webpage content. Such failures may cause incomplete information retrieval and degrade answer quality in affected cases. Second, our experiments primarily focus on English-language benchmarks, and the generalization of GoG to multilingual or cross-lingual visual question answering remains unexplored. Extending the framework to support diverse languages and culturally specific visual knowledge presents an interesting direction for future work.

Acknowledgments

This work is funded in part by the HKUST Start-up Fund (R9911), Theme-based Research Scheme grant (T45-205/21-N), the InnoHK initiative of the Innovation and Technology Commission of the Hong Kong Special Administrative Region Government, and the research funding under HKUST-DXM AI for Finance Joint Laboratory (DXM25EG01).

References

- Sandhini Agarwal, Gretchen Krueger, Jack Clark, Alec Radford, Jong Wook Kim, and Miles Brundage. 2021. Evaluating clip: towards characterization of broader capabilities and downstream implications. *arXiv preprint arXiv:2108.02818*.
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, and 45 others. 2025a. [Qwen3-vl technical report](#). *Preprint*, arXiv:2511.21631.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025b. [Qwen2.5-vl technical report](#). *arXiv preprint arXiv:2502.13923*.
- Yang Chen, Hexiang Hu, Yi Luan, Haitian Sun, Soravit Changpinyo, Alan Ritter, and Ming-Wei Chang. 2023. Can pre-trained vision and language models answer visual information-seeking questions? *arXiv preprint arXiv:2302.11713*.
- Xianfu Cheng, Wei Zhang, Shiwei Zhang, Jian Yang, Xiangyuan Guan, Xianjie Wu, Xiang Li, Ge Zhang, Jiaheng Liu, Yuying Mai, Yutao Zeng, Zhoufutu Wen, Ke Jin, Baorui Wang, Weixiao Zhou, Yunhong Lu, Tongliang Li, Wenhao Huang, and Zhoujun Li. 2025. [Simplevqa: Multimodal factuality evaluation for multimodal large language models](#). *Preprint*, arXiv:2502.13059.
- Mingyang Fu, Yuyang Peng, Dongping Chen, Zetong Zhou, Benlin Liu, Yao Wan, Zhou Zhao, Philip S. Yu, and Ranjay Krishna. 2025. [Seeking and updating with live visual knowledge](#). *Preprint*, arXiv:2504.05288.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Ziniu Hu, Ahmet Iscen, Chen Sun, Zirui Wang, Kai-Wei Chang, Yizhou Sun, Cordelia Schmid, David A Ross, and Alireza Fathi. 2023. Reveal: Retrieval-augmented visual-language pre-training with multi-source multimodal knowledge memory. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 23369–23379.
- Jiaming Ji, Sitong Fang, Wenjing Cao, Jiahao Li, Xuyao Wang, Juntao Dai, Chi-Min Chan, Sirui Han, Yike Guo, and Yaodong Yang. 2025. The mirage of multimodality: Where truth is tested and honesty unravels. *arXiv preprint arXiv:2505.20214*.
- Dongzhi Jiang, Renrui Zhang, Ziyu Guo, Yanmin Wu, Jiayi Lei, Pengshuo Qiu, Pan Lu, Zehui Chen, Chaoyou Fu, Guanglu Song, Peng Gao, Yu Liu, Chunyuan Li, and Hongsheng Li. 2024. [Mmsearch: Benchmarking the potential of large models as multimodal search engines](#). *Preprint*, arXiv:2409.12959.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, and 1 others. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474.
- Geng Li, Jinglin Xu, Yunzhen Zhao, and Yuxin Peng. 2025. Dyfo: A training-free dynamic focus visual search for enhancing llms in fine-grained visual understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 9098–9108.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR.
- Yangning Li, Yinghui Li, Xinyu Wang, Yong Jiang, Zhen Zhang, Xinran Zheng, Hui Wang, Hai-Tao Zheng, Philip S Yu, Fei Huang, and 1 others. 2024. Benchmarking multimodal retrieval augmented generation with dynamic vqa dataset and self-adaptive planning agent. *arXiv preprint arXiv:2411.02937*.

- Yunxin Li, Longyue Wang, Baotian Hu, Xinyu Chen, Wanqi Zhong, Chenyang Lyu, Wei Wang, and Min Zhang. 2023. A comprehensive evaluation of gpt-4v on knowledge-intensive visual question answering. *arXiv preprint arXiv:2311.07536*.
- Weijie Liang, Yuanfeng Song, Xing Chen, Caleb Chen Cao, Sirui Han, and Yike Guo. 2026. Vizomem: A visual-textual memory framework for efficient long-horizon reasoning. In *Findings of the Association for Computational Linguistics: ACL 2026*.
- Sheng-Chieh Lin, Chankyu Lee, Mohammad Shoeybi, Jimmy Lin, Bryan Catanzaro, and Wei Ping. 2024. Mm-embed: Universal multimodal retrieval with multimodal llms. *arXiv preprint arXiv:2411.02571*.
- Weizhe Lin, Jinghong Chen, Jingbiao Mei, Alexandru Coca, and Bill Byrne. 2023. Fine-grained late-interaction multi-modal retrieval for retrieval augmented visual question answering. *Advances in Neural Information Processing Systems*, 36:22820–22840.
- Xiaochuan Liu, Yuanfeng Song, Xiaoming Yin, and Xing Chen. 2026. Datasage: Multi-agent collaboration for insight discovery with external knowledge retrieval, multi-role debating, and multi-path reasoning. In *Findings of the Association for Computational Linguistics: ACL 2026*.
- Junyu Luo, Zhizhuo Kou, Liming Yang, Xiao Luo, Jinsheng Huang, Zhiping Xiao, Jingshu Peng, Chengzhong Liu, Jiaming Ji, Xuanzhe Liu, and 1 others. 2025. Finmme: Benchmark dataset for financial multi-modal reasoning evaluation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 29465–29489.
- Kartik Narayan, Yang Xu, Tian Cao, Kavya Nerella, Vishal M Patel, Navid Shiee, Peter Grasch, Chao Jia, Yinfei Yang, and Zhe Gan. 2025. Deepmmsearch-r1: Empowering multimodal llms in multimodal web search. *arXiv preprint arXiv:2510.12801*.
- OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Madry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, and 401 others. 2024. *Gpt-4o system card*. *Preprint*, arXiv:2410.21276.
- Yayun Qi, Hongxi Li, Yiqi Song, Xinxiao Wu, and Jiebo Luo. 2025. How vision-language tasks benefit from large pre-trained models: A survey. *IEEE Transactions on Multimedia*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. *Gemma 3 technical report*. *Preprint*, arXiv:2503.19786.
- Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. 2024. Eyes wide shut? exploring the visual shortcomings of multimodal llms. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9568–9578.
- Jiaqi Wang, Hanqi Jiang, Yiheng Liu, Chong Ma, Xu Zhang, Yi Pan, Mengyuan Liu, Peiran Gu, Sichen Xia, Wenjun Li, and 1 others. 2024. A comprehensive review of multimodal large language models: Performance and challenges across different tasks. *arXiv preprint arXiv:2408.01319*.
- Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, and 1 others. 2025a. Internvl3. 5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*.
- Xiao Wang, Ibrahim Alabdulmohsin, Daniel Salz, Zhe Li, Keran Rong, and Xiaohua Zhai. 2025b. Scaling pre-training to one hundred billion data for vision language models. *arXiv preprint arXiv:2502.07617*.
- Cong Wei, Yang Chen, Haonan Chen, Hexiang Hu, Ge Zhang, Jie Fu, Alan Ritter, and Wenhua Chen. 2024. Uniir: Training and benchmarking universal multimodal information retrievers. In *European Conference on Computer Vision*, pages 387–404. Springer.
- Jinming Wu, Zihao Deng, Wei Li, Yiding Liu, Bo You, Bo Li, Zejun Ma, and Ziwei Liu. 2025a. Mmsearch-r1: Incentivizing llms to search. *arXiv preprint arXiv:2506.20670*.
- Penghao Wu and Saining Xie. 2024. V?: Guided visual search as a core mechanism in multimodal llms. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13084–13094.
- Yin Wu, Quanyu Long, Jing Li, Jianfei Yu, and Wenya Wang. 2025b. Visual-rag: Benchmarking text-to-image retrieval augmented generation for visual knowledge intensive queries. *arXiv preprint arXiv:2502.16636*.
- Shicheng Xu, Liang Pang, Huawei Shen, Xueqi Cheng, and Tat-Seng Chua. 2024. Search-in-the-chain: Interactively enhancing large language models with search for knowledge-intensive tasks. In *Proceedings of the ACM Web Conference 2024*, pages 1362–1373.

Cui Yakun, Peng Qi, Fushuo Huo, Hang Du, Weijie Shi, Juntao Dai, Zhenghao Zhu, Sirui Han, and Yike Guo. 2026a. *Perception, understanding and reasoning, a multimodal benchmark for video fake news detection*. Preprint, arXiv:2510.24816.

Cui Yakun, Yanting Zhang, Zhu Lei, Jian Xie, Zhizhuo Kou, Hang Du, Zhenghao Zhu, and Sirui Han. 2026b. *Mmfctub: Multi-modal financial credit table understanding benchmark*. arXiv preprint arXiv:2601.04643.

Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Ehsan Azarnasab, Faisal Ahmed, Zicheng Liu, Ce Liu, Michael Zeng, and Lijuan Wang. 2023. *Mmreact: Prompting chatgpt for multimodal reasoning and action*. arXiv preprint arXiv:2303.11381.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2022. *React: Synergizing reasoning and acting in language models*. In *The eleventh international conference on learning representations*.

Michihiro Yasunaga, Armen Aghajanyan, Weijia Shi, Rich James, Jure Leskovec, Percy Liang, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2022. *Retrieval-augmented multimodal language modeling*. arXiv preprint arXiv:2211.12561.

Yujin Zhou, Pengcheng Wen, Jiale Chen, Boqin Yin, Han Zhu, Jiaming Ji, Juntao Dai, Chi-Min Chan, and Sirui Han. 2026. *What, whether and how? unveiling process reward models for thinking with images reasoning*. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 29071–29079.

Han Zhu, Juntao Dai, Jiaming Ji, Haoran Li, Chengkun Cai, Pengcheng Wen, Chi-Min Chan, Boyuan Chen, Yaodong Yang, Sirui Han, and Yike Guo. 2025. *Safemt: Multi-turn safety for multimodal language models*. Preprint, arXiv:2510.12133.

A Ethics Statement

We discuss the ethical considerations and licensing of resources used in this work. Our framework is built upon publicly available models and datasets, and we ensure compliance with their respective licenses.

Model Licenses. We utilize Qwen2.5-VL-7B-Instruct and Qwen3-VL-8B-Think as base models, both released under the Apache-2.0 license, which permits academic research and modification.

Dataset Licenses. All evaluation benchmarks used in this work are publicly available for research purposes:

- **SimpleVQA:** Apache-2.0 license

- **LiveVQA-new:** CC-BY-NC-4.0 license
- **FVQA:** Apache-2.0 license
- **DynVQA:** Apache-2.0 license
- **MMSearch:** Available for research use
- **InfoSeek:** Apache-2.0 license

Human Annotation. The full annotation workflow, including detailed instructions and quality-control procedures, is described in Section 3. All annotators provided informed consent after being briefed on the task objectives, expected workload, and data usage policies. Annotators were compensated at fair market rates for their expertise and time.

Intended Use. Our GoG framework is designed for research purposes in multimodal understanding and knowledge-intensive visual question answering. We do not anticipate direct negative societal impacts from this work. However, as with any system that retrieves information from the web, users should be aware that retrieved content may contain inaccuracies or biases present in online sources.

B Datasets

We evaluate our method on six diverse benchmarks that assess different aspects of knowledge-intensive visual question answering. Below we provide detailed descriptions of each dataset.

B.1 FVQA

FVQA is a factual visual question answering dataset designed to evaluate models on knowledge-requiring visual questions. The dataset is constructed through a combination of automated and manual annotation processes.

The training set, FVQA-train, comprises 5,000 samples with a balanced distribution of approximately 3,400 search-required and 1,600 search-free examples. These samples are collected from three sources: (1) FVQA-auto-vc, containing 5,400 training samples generated by retrieving image-webpage pairs for 10,000 visual concepts sampled from the MetaCLIP metadata distribution, then using GPT-4o to generate factual VQA pairs; (2) FVQA-auto-txt, consisting of 7,000 samples derived from the InfoSeek dataset through balanced sampling across knowledge categories; and (3)

FVQA-manual-train, comprising 800 manually annotated samples where annotators selected knowledge categories, located relevant images, and formulated factual questions with precise answers.

The test set, FVQA-test, contains 1,800 high-quality examples that are either manually verified or fully human-annotated. It includes 600 samples from FVQA-auto-vc with manual verification, 600 samples from the InfoSeek Human Split with manually annotated answers, and 600 samples from direct manual annotation.

B.2 InfoSeek

InfoSeek is a visual question answering dataset tailored for information-seeking questions that cannot be answered with common sense knowledge alone. The dataset is constructed through a semi-automated process that transforms Wikidata triples into natural language questions using human-authored templates. Annotators design question templates for 300 Wikidata relations, incorporating placeholders for visual entity types and units. These questions are paired with corresponding images and answers to form image-question-answer triplets. To ensure diversity and answerability, question-answer pairs lacking supporting evidence in Wikipedia are filtered out, and balanced sampling is applied across entities and relations. We randomly sampled 2,000 examples from its test split for evaluation.

B.3 SimpleVQA

SimpleVQA is a benchmark designed to evaluate factual knowledge boundaries of multimodal large language models. The dataset consists of 2,025 samples spanning 9 core tasks and 9 primary domains. The tasks include Logic & Science, Object Identification Recognition, Time & Event, Person & Emotion, Location & Building, Text Processing, Quantity & Position Relationship, Art & Culture, and Object Attributes Recognition. The domains cover Literature, Education & Sports, Euro-American History & Culture, Contemporary Society, Engineering, Technology & Application, Film, Television & Media, Natural Science, Art, Chinese History & Culture, and Life. All samples follow a short-answer format with standardized answers, enabling objective assessment through direct answer matching.

B.4 MMSearch

MMSearch is a comprehensive evaluation benchmark for assessing multimodal search performance. The dataset contains 300 manually collected instances spanning 14 subfields, with no overlap with current model training data to ensure that correct answers can only be obtained through searching. The benchmark evaluates models on three individual tasks (query, rerank, and summarization) and one end-to-end task involving a complete search process.

B.5 LiveVQA

LiveVQA is a dataset designed to evaluate how multimodal large language models handle up-to-date visual information beyond their training data cutoff. The dataset features 107,143 samples across 12 categories, drawn from recent news articles, video platforms, and academic publications from April 2024 to May 2025. The benchmark specifically tests models on content that extends beyond their knowledge boundaries and evaluates methods for updating models with new visual knowledge.

B.6 Dyn-VQA

Dyn-VQA is a challenging dataset comprising 1,452 dynamic questions that require complex multimodal knowledge retrieval strategies. The dataset includes three types of dynamic questions: (1) questions with rapidly changing answers, where context knowledge updates frequently and retrieved content may mix outdated and newer information; (2) questions requiring multi-modal knowledge, demanding retrieval across diverse modalities with tailored retrieval APIs; and (3) multi-hop questions that necessitate varied reasoning steps for solution. Unlike existing VQA datasets that primarily focus on two-hop questions, Dyn-VQA requires flexible planning of retrieval queries, tools, and timing.

C Tools

This section describes the tools integrated into our framework. We incorporate two categories of tools: search tools (text search and image search) and a grounding tool. The search tools enable the model to retrieve external knowledge from the web, while the grounding tool allows fine-grained visual entity localization before searching.

C.1 Search Tools

Image Search Tool. Our image search tool is built upon SerpAPI. Given an input image URL, SerpAPI returns a set of visually similar webpages along with metadata including URLs, thumbnails, and titles. We rank the returned results by relevance and extract up to five valid results, each represented as a thumbnail-title pair. In our experiments, all input images and grounded visual entities are uploaded to an image hosting service and mapped to corresponding URLs before being passed to the model.

Text Search Tool. The text search tool comprises a complete pipeline consisting of three components: SerpAPI-based text search, Jina Reader, and a summarizer. When the model generates a text search query, it is first sent to SerpAPI, which returns the top-5 retrieval results. Jina Reader then processes the URLs and converts the webpage content into structured text. Finally, we employ Qwen3-32B as a summarizer to extract only the information relevant to the question, filtering out irrelevant content. The entire pipeline is executed in parallel, with Qwen3-32B deployed on $8 \times \text{H800}$ GPUs.

C.2 Grounding Tool

Our grounding tool is based on Grounding DINO. When the model outputs an image search query, the query is passed to our encapsulated Grounding DINO service deployed on $8 \times \text{H800}$ GPUs. Grounding DINO returns the top- n bounding boxes most similar to the query, where $n = 5$ in our implementation. Through manual inspection, we observed that Grounding DINO often returns redundant results with overlapping regions. To address this issue, we propose a Gaze Selection mechanism that selects the 1 to n most question-relevant grounding results. The selected regions are then sent in parallel to the image search tool, and the retrieved information about the grounded content is returned to the model for further reasoning.

D Implementation Details

This section provides comprehensive implementation details for our training and evaluation pipeline, covering three stages: supervised fine-tuning (SFT), reinforcement learning (RL), and evaluation.

Hyperparameter	Value
<i>Supervised Fine-Tuning (SFT)</i>	
Optimizer	AdamW
Learning Rate	1×10^{-5}
LR Scheduler	Cosine
Warmup Ratio	0.1
Batch Size (Global)	128
Epochs	3
Max Sequence Length	32,768
Precision	bf16
DeepSpeed Stage	ZeRO-3
LoRA Rank	8
LoRA Target	All
<i>Reinforcement Learning (RL)</i>	
Algorithm	GRPO
Actor Learning Rate	1×10^{-6}
Critic Learning Rate	1×10^{-5}
KL Penalty in Reward	False
Actor KL Loss Coefficient	0.001
Rollout Sample Size (N)	5
Max Prompt Length	8,192
Max Response Length	8,192
Max Generation Rounds	5
Total Epochs	15
Global Train Batch Size	256
Image Search Limit	3
Text Search Limit	3
Parallel Tool Call Threads	4

Table 6: Hyperparameters for SFT and RL training stages.

D.1 SFT Training Setting

We initialize our model from Qwen2.5-VL-7B-Instruct and Qwen3-VL-Think and perform supervised fine-tuning using the LLaMA-Factory framework. We adopt LoRA (Hu et al., 2022) with rank 8 applied to all target modules for parameter-efficient training. The maximum image resolution is set to 262,144 pixels (equivalent to 512×512) and video resolution to 16,384 pixels. We use the AdamW optimizer with a learning rate of 1×10^{-5} , a cosine learning rate scheduler, and a warmup ratio of 0.1. Training is conducted for 3 epochs with a global batch size of 128, a maximum sequence length of 32,768 tokens, and bf16 mixed precision. We employ DeepSpeed ZeRO-3 for memory-efficient distributed training. The detailed hyperparameters are summarized in Table 6.

D.2 RL Training Setting

For reinforcement learning, we adopt the Group Relative Policy Optimization (GRPO) algorithm within the veRL framework. The actor learning rate is set to 1×10^{-6} with sigmoid decay warmup over 45 steps, while the critic learning rate is 1×10^{-5} .

We disable KL penalty in the reward computation but apply an actor KL loss with coefficient 0.001 to prevent excessive deviation from the reference policy. During rollout, we sample $N = 5$ responses per prompt using vLLM with a maximum prompt length of 8,192 tokens and maximum response length of 8,192 tokens. The model is allowed up to 5 rounds of multi-turn tool interactions per episode. For search constraints, we limit image search and text search to 3 calls each, with a maximum of 5 crop rounds. Tool calls are executed in parallel with 4 threads to accelerate exploration. Training runs for 15 epochs with a global batch size of 256, and checkpoints are saved every 50 steps. The complete hyperparameter configuration is provided in Table 6.

D.3 Evaluation Setting

During inference, we deploy our trained models on $2 \times \text{H800}$ GPUs using vLLM and serve them as online APIs. We use greedy decoding with temperature set to 0 to ensure reproducible results.

For evaluation, we employ an LLM-as-a-Judge approach following prior work (Wu et al., 2025a). Specifically, we use GPT-OSS-120B deployed on $8 \times \text{H800}$ GPUs as the judge model, with temperature set to 0 for deterministic evaluation. The judge model assesses the correctness and quality of generated answers by comparing them against ground-truth references.

E Prompts

This section provides the complete prompts used throughout our framework, including prompts for SFT data generation, inference, summarization, and evaluation. All prompts are carefully designed to guide the model through structured reasoning and tool usage.

E.1 Prompts for SFT Datasets Generation

We design a multi-turn prompt system for SFT data generation that guides the model through different stages of the search process. The system consists of four specialized prompts: (1) **Round 1 Prompt** initializes the reasoning process and instructs the model to analyze the image and select an appropriate action; (2) **After Image Search Prompt** guides the model to process image search results and decide on follow-up actions; (3) **After Gaze Search Prompt** handles the selection of cropped regions for further searching; and (4) **After Text Search Prompt** directs the model to synthesize text search results and formulate the final answer.

Round 1 Prompt

You are an expert visual assistant. Your task is to answer a users question based on the provided image.

Step 1: Analyze the Image

Carefully examine the image and the users question: {question}. Identify all recognizable entities, objects, text, and other visual clues.

Step 2: Plan Your Action

Based on your analysis, you must perform one of the following actions. You must include your thinking process inside a `<think>...</think>` block before choosing an action.

- Action 1: Image Search

If you are not sure about the visual element and need to identify the visual element in the image, you can use one of the following image search methods.

- Whole image search: Only use this if the question is about the entire scene in general, its location, or the overall context. Output only: `<img_search></img_search>`.

Note: Do not output `<img_search></img_search>`.

- Cropped search (recommended for targeted questions): Use this when the question focuses on a particular visual element like an object, person, animal, logo, building, etc. Provide a brief description of the element inside the tags. Examples:

`<img_search>the woman wearing a blue jacket</img_search>`

`<img_search>the emblem on the front of the car</img_search>`

- Action 2: Use Text Search

If you can identify the visual element confidently but need more specific information to answer the question, invoke the text search tool. Generate a focused query and output it as `<text_search>your search query</text_search>`.

- Action 3: Direct Answer

If you are confident in your identification and have enough knowledge to respond, provide a clear and concise answer: `<answer>Your answer here.</answer>`

Remember, search results will be provided to you in subsequent turn. You can analyze the search results and decide your next action. You can perform image search multiple times (e.g., Whole Image Search first, then Cropped Search if needed), and perform multiple text searches to gather relevant information. All search results will be placed inside `<information>...</information>`.

Here is the image and the question:

Question: {question}

Image:

After Image Search Prompt

You have received information from an image search. Your goal is to use this new information to answer the original question: {question}.

Step 1: Analyze the Results

Analyze the information provided within the <information>...</information> block. Consider what you have learned about the visual element in question and how it relates to the user's query.

Step 2: Decide Your Next Action

Include your reasoning inside a <think>...</think> block, then select one of the following actions:

- Action 1: Direct Answer

If the image search results have helped you identify the visual element and you can confidently answer the question with your internal knowledge, provide the final, concise answer inside an <answer>...</answer> tag.

- Action 2: Use Text Search

If the image search results have helped you identify the visual element but you need more specific details to answer the question, invoke the text search tool. Formulate a precise query based on the image search results and output it as <text_search>your search query</text_search>. You can use the text search tool multiple times in subsequent turns if needed.

- Action 3: Search Image Again

If the previous image search results were not helpful (e.g., the crop was inaccurate, or whole image search missed details), you can search the image again.

- If you used Whole Image Search and need details, you MUST try Cropped Search: <img_search>description</img_search>
- If you used Cropped Search and missed, try refining the description: <img_search>refined description</img_search>.

After Gaze Search Prompt

You have requested to crop the image to focus on specific elements. Here are the cropped images based on your request.

Step 1: Analyze the Cropped Images

Review the cropped images provided above. Determine if they clearly show the visual elements you need to answer the question: {question}.

Step 2: Plan Your Next Action

Include your thinking process inside a <think>...</think> block. Then, choose one of the following actions:

● Action 1: Search Selected Crops

If one or more cropped images contain the visual element you need, select them for image search. Output the indices of the crops you want to search (1-based) inside <search_crop>...</search_crop> tags. Example: <search_crop>1,3</search_crop> (to search crop 1 and crop 3).

● Action 2: Answer Directly

If the cropped images themselves provide enough information to answer the question without further searching, provide the final answer inside <answer>...</answer> tags.

● Action 3: Search Again (Crop Again)

If none of the cropped images are satisfactory (e.g., they missed the object or the crop is too loose), you can try to crop again with a different description. Output the new description inside <img_search>...</img_search> tags. Example: <img_search>the license plate on the car</img_search>.

● Action 4: Search Again (Whole Image)

If cropped images are insufficient or the question requires global context that was missed, you can revert to whole image search. Output only: <img_search></img_search>.

Note: At this stage, you MUST NOT use text search (<text_search>). Focus on image-based operations only. Text search will be available after the image search results are returned.

After Text Search Prompt

You have acquired new data from a text search. Your objective is to assess this information and determine the optimal next move to address the original inquiry: {question}.

Step 1: Assess the Findings Examine the content within the <information>...</information> block. Cross-reference it with what you already know to identify any remaining gaps needed to fully answer the question.

Step 2: Determine Strategy First, articulate your reasoning process inside a <think>...</think> block. Then, select one of the following courses of action:

Action 1: Provide Answer If the gathered data is sufficient to cover the request, deliver the final, concise response wrapped in an <answer>...</answer> tag.

Action 2: Refine Search If the current results are useful but incomplete, conduct a follow-up search. Formulate a more targeted or modified query to locate the missing details. Output this new query as <text_search>your refined search query</text_search>.

Action 3: Stop Search If you have exhausted reasonable search options and the answer remains elusive based on available data, conclude by outputting <answer>Unable to answer due to lack of relevant information</answer>.

E.2 Prompts for Direct Answer

To establish a baseline for comparison, we first evaluate models using a direct answer approach where the model receives only the image and question without any external knowledge retrieval. This prompt template instructs the model to provide concise answers based solely on the visual information present in the image.

Direct Answer Prompt

Based on the image, answer the question with as few words as you can.
Question: {question} Image: image

E.3 Prompts for Full Search Workflow

Our full search workflow consists of a two-round prompting strategy designed to leverage external knowledge through simulated search engine interactions.

In the first round, the model is presented with the original question, the input image, and five ranked image search results that provide contextual information related to the query. Each search result includes both a webpage thumbnail and its corresponding title. The model is then tasked with formulating an optimized text query that would effectively retrieve the necessary information from a search engine to answer the original question.

The second round prompt provides the model with the text search results obtained from executing the query generated in the first round. Given these retrieved documents along with the original question and image, the model is instructed to synthesize the information and produce a concise final answer.

1st Round Prompt

You are presented with a question and an image that requires external knowledge to be answered. To aid your understanding, five ranked image search results related to the original image are provided. Each result contains a webpage image and its title.

Question: {question} Image: {image}

Image Search Results:

1. Webpage Image: {image} Webpage Title: {title}
2. Webpage Image: {image} Webpage Title: {title}
- ...
5. Webpage Image: {image} Webpage Title: {title}

Assume you have access to a search engine (e.g., Google). Based on the question, the original image, and the search results, formulate a targeted text query to retrieve information necessary to answer the question correctly. Optimize your query for search engine algorithms.

Output the text query directly (strictly without quotes or explanation):

2nt Round Prompt

Please read the text search results and answer the question based on the provided image.
Text Search Results: ...
Original question: {question}
Keep your answer as concise as possible.

E.4 Prompts for Summarization

After retrieving information from text search, we use a summarization prompt to extract only the relevant content from the retrieved webpages. This step is crucial for filtering out noise and reducing the context length before feeding the information back to the model.

Summarization Prompt

Synthesize the text extracted from Google Lens search results (titles and descriptions) to describe the primary visual content of the original image. Focus on identifying key elements such as figures, objects, settings, or text. Limit your response to 4-5 sentences.
Extracted Text:
{formatted_results}

E.5 Prompts for LLM as a Judge

We adopt an LLM-as-a-Judge approach for evaluation, where a separate language model assesses whether the predicted answer is semantically equivalent to the ground-truth answer. The judge is instructed to consider alternate correct answers and focus on semantic equivalence rather than exact string matching.

LLM as a Judge Prompt

You are an impartial judge evaluating a model's answer for a visual question answering task.

Question: {question}
Ground-Truth Answer(s): {gt_str}
Predicted Answer: {prediction}

IMPORTANT:
- The Ground-Truth Answer(s) may contain alternate correct answers
- The predicted answer is CORRECT if it is semantically equivalent to at least ONE ground-truth answer
- Respond with ONLY one of these two options:
[CORRECT]
[INCORRECT]

Your response:

E.6 Prompts for Reward Model

During reinforcement learning, we use an LLM as reward model to evaluate the quality of generated responses.

LLM as a Reward Model Prompt

You are an intelligent and impartial expert evaluator for Question Answering tasks. Your specific task is to determine whether the "Predicted Answer" is factually consistent with the "Ground Truth Answer(s)".

Input Data

- **Question**: {question}
- **Ground Truth Answer(s)**: {ground_truth} (A list of one or more acceptable reference answers).
- **Predicted Answer**: {prediction} (The model's response to evaluate).

Evaluation Criteria

Please assess the alignment using the following comprehensive rules. If the Predicted Answer matches **ANY** of the Ground Truth Answers according to these rules, the final result is **True**.

- Semantic Equivalence (Core Meaning)**
 - **Synonyms & Aliases**: Accept common abbreviations (e.g., "USA" = "United States"), acronyms, and synonyms.
 - **Formatting**: Ignore differences in casing (lower/upper), punctuation, articles ("a", "the"), and spacing.
 - **Typos**: Ignore minor spelling errors if the core meaning is unambiguous.
- Numerical & Logical Flexibility**
 - **Units**: Accept equivalent values in different units (e.g., "1.5 kg" = "1500 g", "\$1M" = "1 million dollars").
 - **Rounding**: Accept reasonable rounding deviations (e.g., "3.14" matches "3.14159").
 - **Ranges**: If the prediction is a range that **contains** the specific ground truth value, it is a MATCH (e.g., Prediction "1990-2000" matches Truth "1995").
 - **Dates**: Accept different date formats (e.g., "Jan 1, 2020" = "2020/01/01").
- Specificity & Information Containment**
 - **Superset (Extra Info)**: If the prediction contains the correct answer **plus** additional context that does **NOT** contradict the facts, it is CORRECT. (e.g., Truth: "Paris", Prediction: "Paris, the capital of France" -> MATCH).
 - **Subset (More Specific)**: If the prediction is a more precise/specific version of a general ground truth, it is CORRECT.
 - **Lists**: If the answer is a list of items, the order does NOT matter (e.g., "A and B" = "B and A").
- Negative Constraints (When to Reject)**
 - **Contradiction**: If the prediction includes the correct entity but makes a claim that directly **contradicts** the question or truth, it is FALSE.
 - **Wrong Entity**: If the core entity is different (e.g., "Mars" vs "Jupiter"), strict rejection.
 - **Hallucination**: If the prediction adds extra details that are factually wrong and change the answer's validity, it is FALSE.

Output Format

You must output your decision strictly in the following XML format:

```
<reason>
[Step-by-step reasoning: 1. Identify key entities in Truth and Prediction. 2. Check for semantic/numerical match. 3. Check
for contradictions. Explain why it matches or fails.]
</reason>
<judge>
True
</judge>
```

(Output "True" if it matches, "False" otherwise)